arXiv:1601.07600v2 [math.OC] 30 Jan 2016

SHRINAKGE FUNCTION AND ITS APPLICATIONS IN MATRIX APPROXIMATION

TOBY BOAS*, ARITRA DUTTA†, XIN LI‡, KATIE MERCIER§, AND ERIC NIDERMAN¶

Abstract. We first give an elementary derivation of the shrinkage function widely used in compressive sensing and statistical estimations. Then we demonstrate how it can be used in solving several well-known problems and proving one new result in matrix approximations.

Key words. Shrinkage function, singular value decomposition, low-rank approximation, sparse approximation.

AMS subject classifications. 65F15, 65F30, 65F35, 65F50, 65K10

1. Introduction. There are a few mathematical functions which are introduced for convenience. For example, the Heaviside step function $H(\cdot)$, a piecewise constant function given by:

$$H(x) = \begin{cases} 1, & x > 0 \\ \frac{1}{2}, & x = 0 \\ 0, & x < 0 \end{cases},$$

and the Dirac delta function $\delta(\cdot)$ (more precisely, a distribution (see, e.g., [1])), a generalized function whose discrete analog is referred to as the Kronecker delta function:

$$\delta_{ij} = \begin{cases} 1, & i = j \\ 0, & i \neq j \end{cases}.$$

This article is on a newcomer, the shrinkage function $S_\lambda(\cdot)$, first introduced by Donoho and Johnstone in their landmark paper ([2], see also [3]) on function estimation using wavelets in the early 1990's. Recently, the shrinkage function has been heavily used in the solutions of several optimization and approximation problems of matrices (see, e.g., [4, 5, 6, 7]). We give an elementary treatment that is accessible to a vast group of researchers, as it only requires basic knowledge in calculus and linear algebra and show how naturally the shrinkage function can be used in solving more advanced problems.

2. A calculus problem. We start with a regular calculus problem. Let $\lambda > 0$ and $a \in \mathbb{R}$ be given. Consider the following problem:

$$\min_{x \in \mathbb{R}} \left[\lambda |x| + \frac{1}{2}(x-a)^2 \right]. \tag{2.1}$$

*Department of Mathematics, University of Florida, Gainesville, FL 32611-8105 (tboas@ufl.edu).
†Department of Mathematics, University of Central Florida, 4000 Central Florida Blvd, Orlando, Florida 32816 (d.aritra2010@knights.ucf.edu).
‡Department of Mathematics, University of Central Florida, 4000 Central Florida Blvd, Orlando, Florida 32816 (xli@ucf.edu).
§656 Eldron Ave, Deltona, FL, 32738(katie.mercier@knights.ucf.edu).
¶Department of Mathematics, University of Central Florida, 4000 Central Florida Blvd, Orlando, Florida 32816 (eaniederman@knights.ucf.edu).

We adopt the notation $a = \arg\min_{x\in A} f(x)$ to mean that $a \in A$ is a solution of the minimization problem $\min_{x\in A} f(x)$ and define:

$$S_\lambda(a) := \arg\min_{x\in\mathbb{R}} \left[\lambda|x| + \frac{1}{2}(x-a)^2\right]. \tag{2.2}$$

THEOREM 2.1. *Let $\lambda > 0$ be fixed. For each $a \in \mathbb{R}$, there is one and only one solution $S_\lambda(a)$, to the minimization problem (2.2). Furthermore,*

$$S_\lambda(a) = \begin{cases} a-\lambda, & a > \lambda \\ 0, & |a| \le \lambda \\ a+\lambda, & a < -\lambda \end{cases}.$$

Remark. The function $S_\lambda(\cdot)$ defined above is called the shrinkage function (also referred to as soft shrinkage or soft threshold, [2, 3]). One may imagine that $S_\lambda(a)$ "shrinks" a to zero when $|a| \le \lambda$. A plot of $S_\lambda(\cdot)$ for $\lambda = 1$ is given in Figure 2.1.

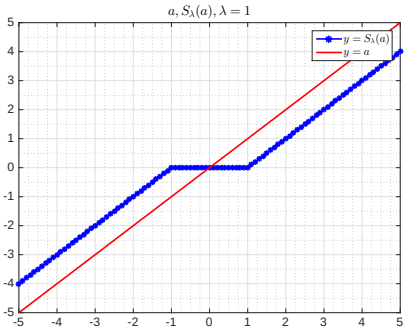

Fig. 2.1: A plot of S_λ for $\lambda = 1$.

Proof. Let $f(x) = \lambda|x| + \frac{1}{2}(x-a)^2$. Note that $f(x) \to \infty$ when $|x| \to \infty$ and f is continuous on $\mathbb{R}$ and differentiable everywhere except a single point $x = 0$. So, f achieves its minimum value on $\mathbb{R}$ at one of its critical points. Let $x^* = \arg\min_{x\in\mathbb{R}} f(x)$.

We consider three cases.

Case 1: $x^* > 0$. Since f is differentiable at $x = x^*$ and achieves its minimum, we must have $f'(x^*) = 0$. Note that, for $x > 0$, we have

$$f'(x) = \frac{d}{dx}[\lambda x + \frac{1}{2}(x-a)^2] = \lambda + (x-a).$$

So,

$$\lambda + (x^* - a) = 0,$$

which implies

$$x^* = a - \lambda.$$

To be consistent with $x^* > 0$, it is necessary that $a - \lambda > 0$ or, equivalently, $a > \lambda$.

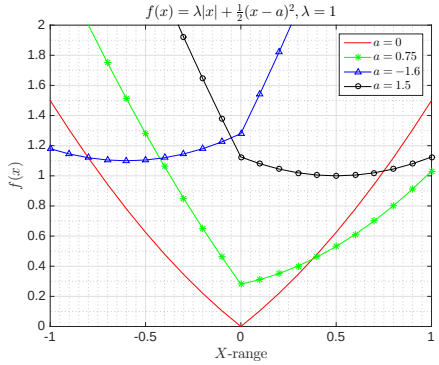

Fig. 2.2: Plots of $f(x)$ for different values of a with $\lambda = 1$.

Case 2: $x^* < 0$. By proceeding similarly as in Case 1 above, we can arrive at

$$x^* = a + \lambda \quad \text{with} \quad a < -\lambda.$$

Case 3: $x^* = 0$. Note that $f(x)$ is no longer differentiable at $x^* = 0$ (so we could not use the condition $f'(x^*) = 0$ as before). But since f has a minimum at $x^* = 0$ and since f is differentiable on each side of $x^* = 0$, it is necessary that

$$f'(x) > 0 \ \text{ for } \ x > 0 \ \text{ and } \ f'(x) < 0 \text{ for } \ x < 0.$$

So,

$$\lambda + x - a > 0 \ \text{ for } x > 0 \ \text{ and } \ -\lambda + x - a < 0 \ \text{ for } x < 0.$$

Thus,

$$\lambda - a > 0 \ \text{ and } \ -\lambda - a < 0,$$

or, equivalently,

$$|a| \leq \lambda.$$

To summarize, we have

$$x^* = \begin{cases} a - \lambda & \text{with } a > \lambda, \\ a + \lambda & \text{with } a < -\lambda, \\ 0 & \text{with } |a| \leq \lambda. \end{cases}$$

Since one and only one of the three cases (1) $a > \lambda$, (2) $a < -\lambda$, and (3) $|a| \leq \lambda$ holds, we obtain the uniqueness in general. With the uniqueness, it is straightforward to verify that each of the three cases would imply the corresponding formula for x^*. This completes the proof. □

3. A sparse recovery problem. Recently, research in compressive sensing leads to the recognition of the fact that the ℓ_1-norm of a vector is a good substitute for the count of the number of non-zero entries of the vector in many minimization problems. In this section, we solve some simple minimization problems using the count of non-zero entries or the ℓ_1-norm. Given a vector $\mathbf{v} \in \mathbb{R}^n$, we want to solve

$$\min_{\mathbf{u}\in\mathbb{R}^n} [\text{card}(\mathbf{u}) + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2], \tag{3.1}$$

where $\text{card}(\mathbf{u})$ denotes the number of non-zero entries of $\mathbf{u}$, $\|\cdot\|_{\ell_2}$ denotes the Euclidean norm in $\mathbb{R}^n$, and $\beta > 0$ is a given balancing parameter. We can solve problem (3.1) component-wise (in each u_i) as follows. Notice that, given $\mathbf{u} \in \mathbb{R}^n$, each entry u_i of $\mathbf{u}$ contributes 1 to $\text{card}(\mathbf{u})$ if u_i is non-zero, and contributes 0 if u_i is zero. Since we are minimizing $g(\mathbf{u}) := \text{card}(\mathbf{u}) + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2$, if u_i is zero then the contribution to $g(\mathbf{u})$ depending on this u_i is $\frac{\beta}{2}v_i^2$; otherwise, if u_i is non-zero, then we should minimize $\frac{\beta}{2}(u_i - v_i)^2$ for $u_i \in \mathbb{R}\setminus\{0\}$, which forces that $u_i = v_i$ and contributes 1 to $g(\mathbf{u})$ as the minimum value. Therefore, the solution $\mathbf{u}$ to problem (3.1) is given component-wise by

$$u_i = \begin{cases} 0, & \text{if } \frac{\beta}{2}(v_i)^2 \le 1 \\ v_i, & \text{otherwise.} \end{cases}$$

Next, we replace $card(\mathbf{u})$ by $\|\mathbf{u}\|_{l_1}$ in (3.1) and solve:

$$\min_{\mathbf{u}\in\mathbb{R}^n} [\|\mathbf{u}\|_{\ell_1} + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2], \tag{3.2}$$

where $\|\cdot\|_{\ell_1}$ denotes the ℓ_1 norm in $\mathbb{R}^n$.

Using Theorem 2.1, we can solve (3.2) component-wise as follows.

THEOREM 3.1. *[8] Let $\beta > 0$ and $\mathbf{v} \in \mathbb{R}^n$ be given and let*

$$\mathbf{u}^* = \arg\min_{\mathbf{u}\in\mathbb{R}^n} [\|\mathbf{u}\|_{\ell_1} + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2],$$

then

$$\mathbf{u}^* = S_{1/\beta}(\mathbf{v}),$$

where, $S_{1/\beta}(\mathbf{v})$ denotes the vector whose entries are obtained by applying the shrinkage function $S_{1/\beta}(\cdot)$ to the corresponding entries of $\mathbf{v}$.

Proof. If u_i and v_i denote the ith entry of the vectors $\mathbf{u}$ and $\mathbf{v}$, respectively, $i = 1, 2, \ldots, n$, then we have,

$$\begin{aligned} \mathbf{u}^* &= \arg\min_{\mathbf{u}\in\mathbb{R}^n} [\|\mathbf{u}\|_{\ell_1} + \frac{\beta}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2] \\ &= \arg\min_{\mathbf{u}\in\mathbb{R}^n} \left[\sum_{i=1}^{n} \left(\frac{1}{\beta}|u_i| + \frac{1}{2}(u_i - v_i)^2\right)\right]. \end{aligned}$$

Since the i-th term in the summation depends only on u_i, the vector $\mathbf{u}^*$ must have components u_i^* satisfying

$$u_i^* = \arg\min_{u_i^*\in\mathbb{R}} [\frac{1}{\beta}|u_i| + \frac{1}{2}(u_i - v_i)^2],$$

for $i = 1, 2, \ldots, n$. But by Theorem 2.1, the solution to each of these problems is given precisely by $S_{1/\beta}(v_i)$. This yields the result. □

Remark. The previous proof still works if we replace the vectors by matrices and use the extension of the ℓ_1 and ℓ_2-norms to matrices *by treating them as vectors*. Thus by using the same argument we can easily show the following matrix version of the previous theorem.

THEOREM 3.2. *[8] Let $\beta > 0$ and $\mathbf{V} \in \mathbb{R}^{m\times n}$ be given. Then*

$$S_{1/\beta}(\mathbf{V}) = \arg\min_{\mathbf{U}\in\mathbb{R}^{m\times n}} [\|\mathbf{U}\|_{\ell_1} + \frac{\beta}{2}\|\mathbf{U} - \mathbf{V}\|_{\ell_2}^2],$$

where $S_{1/\beta}(\mathbf{V})$ is again defined component-wise.

Theorem 3.2 solves the problem of approximating a given matrix by a sparse matrix by using the shrinkage function.

4. Approximation by low rank matrices. The sparse approximation as given by Theorem 3.2 has many applications such as data compression and dimension reduction. In these areas, one may also be interested in finding matrices of low rank. For example, given a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, we want to solve the following approximation problem:

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \left[\text{rank}(\mathbf{X}) + \frac{\beta}{2}\|\mathbf{X} - \mathbf{A}\|_F^2\right],$$

where $\|\cdot\|_F$ denotes the Frobenius norm of matrices (which turns out to be equivalent to the vector l_2 norm if we treat a matrix as a vector - see more discussion on the matrix norms in subsection 4.1).

This is a harder problem since rank(X) is not a convex function. A convex relaxation (see, e.g., [9]) of the problem is provided by replacing the term rank$(\mathbf{X})$ by the nuclear norm of $\mathbf{X}$, $\|\mathbf{X}\|_*$, (again, see subsection 4.1 for a discussion on the nuclear norm and its properties). The problem then becomes:

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \left[\|\mathbf{X}\|_* + \frac{\beta}{2}\|\mathbf{X} - \mathbf{A}\|_F^2\right]. \tag{4.1}$$

This problem again yields an explicit solution ([4, 10]), but in these literatures, the formula is derived by using advanced tools from convex analysis ("subdifferentials" to be more specific). Here, we will show how we can obtain the solution by using simple ideas from the previous section.

4.1. Singular value decomposition and matrix norms. It will be beneficial to recall the various matrix norms. Many useful matrix norms can be defined in terms of the singular values of the matrices. We will deal with two of them: the nuclear norm $\|\cdot\|_*$ and the Frobenius norm $\|\cdot\|_F$.

Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ and $\mathbf{A} = \mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ be a singular value decomposition (SVD) of $\mathbf{A}$ with $\mathbf{U} \in \mathbb{R}^{m\times m}$ and $\mathbf{V} \in \mathbb{R}^{n\times n}$ being two orthogonal matrices (that is, $\mathbf{U}^{-1} = \mathbf{U}^T$ and $\mathbf{V}^{-1} = \mathbf{V}^T$) and $\tilde{\mathbf{A}} = \text{diag}(\sigma_1(\mathbf{A})\ \sigma_2(\mathbf{A})\cdots\sigma_{\min\{m,n\}}(\mathbf{A}))$ being a diagonal matrix such that $\sigma_1(\mathbf{A}) \geq \sigma_2(\mathbf{A}) \geq \cdots \geq \sigma_{\min\{m,n\}}(\mathbf{A}) \geq 0$. The $\sigma_i(\mathbf{A})$'s are called the singular values of $\mathbf{A}$. It is known ([11]) that every matrix in $\mathbb{R}^{m\times n}$ has a SVD and that SVD of a matrix is not unique. Then the nuclear norm of $\mathbf{A}$ is given by

$$\|\mathbf{A}\|_* = \sum_{i=1}^{\min\{m,n\}} \sigma_i(\mathbf{A}),$$

and we can also define the Frobenius norm of $\mathbf{A}$ as

$$\|\mathbf{A}\|_F = \left(\sum_{i=1}^{\min\{m,n\}} (\sigma_i(\mathbf{A}))^2 \right)^{1/2}.$$

This norm turns out to be the same as the ℓ_2 norm of $\mathbf{A}$, treated as a vector in $\mathbb{R}^{mn\times 1}$. This is because the nonzero singular values $\sigma_i(\mathbf{A})$'s are exactly the square root of the nonzero eigenvalues of $\mathbf{A}\mathbf{A}^T$ or $\mathbf{A}^T\mathbf{A}$. So,

$$\|\mathbf{A}\|_{l_2}^2 = \sum_{i=1}^{m}\sum_{j=1}^{n}(a_{ij})^2 = \operatorname{trace}(\mathbf{A}\mathbf{A}^T) = \sum_{i=1}^{\min\{m,n\}} (\sigma_i(\mathbf{A}))^2.$$

Here we have used trace$(\cdot)$ to denote the trace of a matrix (which is equal to the sum of all diagonal entries of the matrix).

We will need the following simple fact about the nuclear norms of a matrix and that of its diagonal: Let $D(\mathbf{A})$ denote the diagonal matrix using the diagonal of $\mathbf{A}$. We have

$$\|D(\mathbf{A})\|_* \le \|\mathbf{A}\|_*. \tag{4.2}$$

This inequality can be verified by using a SVD of $\mathbf{A} = \mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ as follows. Write $\mathbf{U} = (u_{ij})$, $\mathbf{V} = (v_{ij})$, and $t = \min\{m, n\}$. Then

$$\|D(\mathbf{A})\|_* = \|D(\mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T)\|_* = \sum_{i=1}^{m}\left|\sum_{j=1}^{t}\sigma_j(\mathbf{A})u_{ij}v_{ij}\right| \le \sum_{j=1}^{t}\sigma_j(\mathbf{A})\sum_{i=1}^{m}|u_{ij}v_{ij}|$$

$$\le \sum_{j=1}^{t}\sigma_j(\mathbf{A})\cdot\left(\sum_{i=1}^{m}|u_{ij}|^2\right)^{1/2}\left(\sum_{i=1}^{m}|v_{ij}|^2\right)^{1/2} \le \sum_{j=1}^{t}\sigma_j(\mathbf{A}) = \|\mathbf{A}\|_*,$$

where we have used the Cauchy-Schwarz inequality in obtaining the second inequality, and the orthogonality of $\mathbf{U}$ and $\mathbf{V}$ (so that $\sum_{i=1}^{m}|u_{ij}|^2 \le 1$ and $\sum_{i=1}^{m}|v_{ij}|^2 \le 1$) in the last inequality.

We will also use the fact that for any orthogonal matrices $\mathbf{O}$ and $\mathbf{R}$, $\mathbf{OAR}$ and $\mathbf{A}$ have the same singular values, and therefore their Frobenius norms and nuclear norms are same:

$$\|\mathbf{OAR}\|_F = \|\mathbf{A}\|_F \text{ and } \|\mathbf{OAR}\|_* = \|\mathbf{A}\|_*.$$

This is known as the unitary invariance of the Frobenius norm and nuclear norm.

4.2. Solution to (4.1) via problem (2.2). We are ready to show how problem (4.1) is problem (2.2) in disguise. Given $\beta > 0$, using the unitary invariance of the Frobenius norm and the nuclear norm, we have

$$\begin{aligned}\min_{\mathbf{X}}[\|\mathbf{X}\|_* + \frac{\beta}{2}\|\mathbf{X}-\mathbf{A}\|_F^2] &= \min_{\mathbf{X}}[\|\mathbf{X}\|_* + \frac{\beta}{2}\|\mathbf{X}-\mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T\|_F^2] \\ &= \min_{\mathbf{X}}[\|\mathbf{U}^T\mathbf{X}\mathbf{V}\|_* + \frac{\beta}{2}\|\mathbf{U}^T\mathbf{X}\mathbf{V}-\tilde{\mathbf{A}}\|_F^2].\end{aligned}$$

It can be seen from the last expression that the minimum occurs when $\tilde{\mathbf{X}} := \mathbf{U}^T\mathbf{X}\mathbf{V}$ is diagonal since both terms in that expression get no larger when $\tilde{\mathbf{X}}$ is replaced by its diagonal matrix (with the help of (4.2)). So, the matrix $\mathbf{E} = (e_{ij}) := \tilde{\mathbf{X}} - \tilde{\mathbf{A}}$ is a diagonal matrix. Thus,

$$\mathbf{X} = \mathbf{U}\tilde{\mathbf{X}}\mathbf{V}^T, \text{ with } \tilde{\mathbf{X}} = \tilde{\mathbf{A}} + \mathbf{E},$$

which yields a SVD of $\mathbf{X}$ (using the same matrices $\mathbf{U}$ and $\mathbf{V}$ as in a SVD of $\mathbf{A}$!). Then,

$$\begin{aligned}\min_{\mathbf{X}}[\|\mathbf{X}\|_* + \frac{\beta}{2}\|\mathbf{X}-\mathbf{A}\|_F^2] &= \min_{\tilde{\mathbf{X}}\in\text{diag}}\ [\|\tilde{\mathbf{X}}\|_* + \frac{\beta}{2}\|\tilde{\mathbf{X}}-\tilde{\mathbf{A}}\|_F^2] \\ &= \min_{\tilde{\mathbf{X}}\in\text{diag}}\ [\sum_i \tilde{x}_{ii} + \frac{\beta}{2}\sum_i(\tilde{x}_{ii}-\sigma_i(\mathbf{A}))^2],\end{aligned}$$

where "diag" is the set of diagonal matrices in $\mathbb{R}^{m\times n}$. Above is an optimization problem like (2.2) (for vectors $(\tilde{x}_{11}, \tilde{x}_{22}, \cdots)^T$ as $\tilde{\mathbf{X}}$ varies) whose solution is given by

$$\tilde{x}_{ii} = S_{1/\beta}(\sigma_i(\mathbf{A})),\ i = 1, 2, \cdots .$$

To summarize, we have proven the following.

THEOREM 4.1. *[4] Suppose that $\mathbf{A} \in \mathbb{R}^{m\times n}$ and $\beta > 0$ are given. Then the solution to the minimization problem (4.1) is given by $\hat{\mathbf{X}} = \mathbf{U}\tilde{\mathbf{X}}\mathbf{V}^T$ where the diagonal matrix $\tilde{\mathbf{X}}$ has diagonal entries*

$$\tilde{x}_{ii} = S_{1/\beta}(\sigma_i(\mathbf{A})),\ i = 1, 2, \ldots, \min\{m, n\},$$

where $\mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ is a SVD of $\mathbf{A}$.

Remark. 1. A recent proof of this theorem is given by Cai, Candes, and Shen in [4] where they give an advanced verification of the result. Our proof given above has the advantage that it is elementary and allows the reader to "discover" the result.

2. There are many earlier discoveries of related results ([12]) where rank($\mathbf{X}$) is used instead of the nuclear norm $\|\mathbf{X}\|_*$. We will examine one such variant in the next section.

3. One key ingredient in the above discussion is the unitary invariance of the norms $\|\cdot\|_*$ and $\|\cdot\|_F$. It was von Neumann (see, e.g., [9]) who was among the first to study the family of all unitarily invariant matrix norms in matrix approximation, $\|\cdot\|_F$ being one of them.

4. A closely related (but harder) problem is *compressive sensing* ([13, 14]). Readers are strongly recommended to the recently survey by Bryan and Leise ([15]).

5. A variation. More problems can be solved by applying similar ideas. For example, let us consider a variant of a well-known result of Schmidt (see, e.g., [12, Section 5]), replacing the rank by the nuclear norm: For a fixed positive number τ, consider

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{X}-\mathbf{A}\|_F \ \text{ subject to } \ \|\mathbf{X}\|_* \le \tau. \tag{5.1}$$

Using similar methods as in Section 3, this problem can be transformed into the following:

$$\min_{\mathbf{u}\in\mathbb{R}^{\min\{m,n\}}} \|\mathbf{u}-\mathbf{v}\|_{\ell_2} \ \text{ subject to } \ \|\mathbf{u}\|_{\ell_1} \le \tau. \tag{5.2}$$

Note that, (5.2) is related to a Lasso problem [3, 16, 17]. But unlike a LASSO problem, no special assumption (like 0 mean) is made on $\mathbf{v}$ in (5.2). As in [3], one can form a *Lagrangian* of (5.2) and solve:

$$\mathbf{u}^* = \arg \min_{\mathbf{u}\in\mathbb{R}^{\min\{m,n\}}} \{\frac{1}{2}\|\mathbf{u}-\mathbf{v}\|_{\ell_2}^2 + \lambda\|\mathbf{u}\|_{\ell_1}\}, \text{ with } \|S_\lambda(\mathbf{v})\|_{\ell_1} = \tau, \tag{5.3}$$

which has a solution $\mathbf{u}^* = S_\lambda(\mathbf{v})$ according to Theorem 3.1. (The reason for us to use λ instead of β in (5.2) is nonessential: it is only for indicating the similarity with LASSO formulation.) We now sketch the derivation of converting (5.1) to (5.2): As before, let $\mathbf{A} = \mathbf{U}\tilde{\mathbf{A}}\mathbf{V}^T$ be a SVD of $\mathbf{A}$. Then,

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{X}-\mathbf{A}\|_F = \min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{U^T X V} - \tilde{\mathbf{A}}\|_F.$$

Note that $\|\mathbf{X}\|_* = \|\mathbf{U}^T\mathbf{X}\mathbf{V}\|_*$, so (5.1) can be written as

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{X}-\tilde{\mathbf{A}}\|_F \text{ subject to } \|\mathbf{X}\|_* \le \tau,$$

which, by using (4.2), can be further transformed to

$$\min_{\mathbf{X}\in\mathbb{R}^{m\times n}} \|\mathbf{X}-\tilde{\mathbf{A}}\|_F \text{ subject to } \mathbf{X} \text{ being diagonal and } \|\mathbf{X}\|_* \le \tau. \tag{5.4}$$

Next, if we let $\mathbf{u}$ and $\mathbf{v}$ be two vectors in $\mathbb{R}^{\min\{m,n\}}$ consisting of the diagonal elements of $\mathbf{X}$ and $\tilde{\mathbf{A}}$, respectively, then (5.4) is (5.2). Thus we have established the following result.

THEOREM 5.1. *With the notations above, the solution to problem (5.1) is given by*

$$\hat{\mathbf{X}} = \mathbf{U}S_\lambda(\tilde{\mathbf{A}})\mathbf{V}^T,$$

for some λ such that $\|S_\lambda(\tilde{\mathbf{A}})\|_{\ell_1} = \tau$.

Acknowledgment. This work is partially supported by a CSUMS program of the National Science Foundation DMS-0803059. Toby Boas, Katie Mercier, and Eric Niederman contributed to this work as undergraduate students.

REFERENCES

[1] R. BRACEWELL, The Impulse Symbol, *The Fourier Transform and Its Applications* (3rd ed., 2000), Ch. 5, McGraw-Hill, 74–104.
[2] D. L. DONOHO AND I. M. JOHNSTONE, *Ideal spatial adaptation by wavelet shrinkage*, Biometrika, 81 (1994) 425–455.
[3] R. TIBSHIRANI, *Regression shrinkage and selection via the LASSO*, J. of the Royal statistical society, series B, 58 (1996) 267–288.
[4] J. CAI, E. J. CANDES, AND Z. SHEN, *A singular value thresholding algorithm for matrix completion*, SIAM J. Optim., 20 (2010) 1956–1982.
[5] Z. LIN, M. CHEN, AND Y. MA, *The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices*, arXiv:1009.5055v2, March 9, 2011.
[6] M. TAO AND X. YUAN, *Recovering low-rank and sparse components of matrices from incomplete and noisy observations*, SIAM J. Optim., 21 (2011) 57–81.
[7] X. YUAN AND J. YANG, *Sparse and low-rank matrix decomposition via alternating direction methods*, Technical report (available from http://www.optimization-online.org/DBFILE/2009/11/2447.pdf), Dept. of Mathematics, Hong Kong Baptist University, 2009.

[8] W. Yin, E. Hale, and Y. Zhang, *Fixed-point continuation for l_1-minimization: methodology and convergence*, SIAM J. Optim., 19 (2008) 1107–1130.
[9] M. Fazel, *Matrix Rank Minimization with Applications*, Ph.D. dissertation, Department of Electrical Engineering, Stanford University, 2002.
[10] S. Q. Ma, D. Goldfarb and L. F. Chen, *Fixed point and Bregman iterative methods for matrix rank minimization*, Math. Prog. Ser. A, (2009) DOI 10.1007/s10107-009-0306-5.
[11] G. W. Strang, *Introduction to Linear Algebra*, 3rd ed., Wellesley-Cambridge Press, 1998.
[12] G. W. Stewart, *On the early history of the singular value decomposition*, SIAM Review, 35 (1993) 551–566.
[13] E. J. Candes, J. Romberg, and T. Tao, *Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information*, IEEE Transactions on Information Theory, 52 (2006) 489–509.
[14] E. J. Candes and T. Tao, *The power of convex relaxation: Near-optimal matrix completion*, IEEE Transactions on Information Theory, 56 (2010) 2053–2080.
[15] K. Bryan and T. Leise, *Making do with less: an introduction to compressed sensing*, SIAM Review, 55 (2013), 547–566.
[16] R. J. Tibshirani and J. Taylor, *The solution path of the generalized LASSO*, The Annals of Statistics, 39-3 (2011), 1335–1371.
[17] M. R. Osborne, B. Presnell, and B. A. Turlach, *On the LASSO and its dual*, J. of Computational and Graphical Stat, 9 (1999), 319–337.